The term structure of carbon premia

ABFER 10th Annual Conference, Singapore, May 2023

Dora Xia and Omar Zulaica

Bank for International Settlements

Motivation

- It is an emerging consensus that transitioning the global economy to a low-carbon growth path is essential.
  - GHGs increases ➔ global temperature rises ➔ more frequent and intensive nature disasters

- Pricing this transition, investors will likely demand compensation for investing firms with higher carbon footprints, giving rise to ***a carbon premium***. Two channels could be at work, including
  i. The **preference channel**, where investors might have a preference, all else equal, for firms that they perceive as helping to achieve sustainable growth.
  ii. The **risk channel**, where investors perceive carbon-intensive firms as more prone to default.
  iii. The two channels corresponds to investment processes such as negative screening and rating agency's practice of considering environmental factors in assigning risk grades respectively.

- Why is it important to analyze carbon premia?
  - A precondition for capital market to facility transition to sustainable growth is that investors differentiate between financial assets that fund activities with different degrees of environmental impact.
  - Whether climate risk is sufficiently priced in financial assets has important financial stability implications.

Our paper in a nutshell

- What do we do
  - Whether such carbon premium exists in corporate bonds
  - Why corporate bonds
    - Represent an important source for corporates to obtain funding
    - Offer a good foundation to dissect the preference and risk angles
    - Complement other work on carbon premium in equity and bank lending

- How do we do it
  - Illustration

- What do we find
  - Both the preference and risk channels are priced in.
  - Carbon premium is larger for bonds issued by firms in energy-intensive sectors.
  - Carbon premium loads differently across maturities – a hump-shaped term structure.

Our paper in an illustration

Determinants of corporate bond spreads: illustration

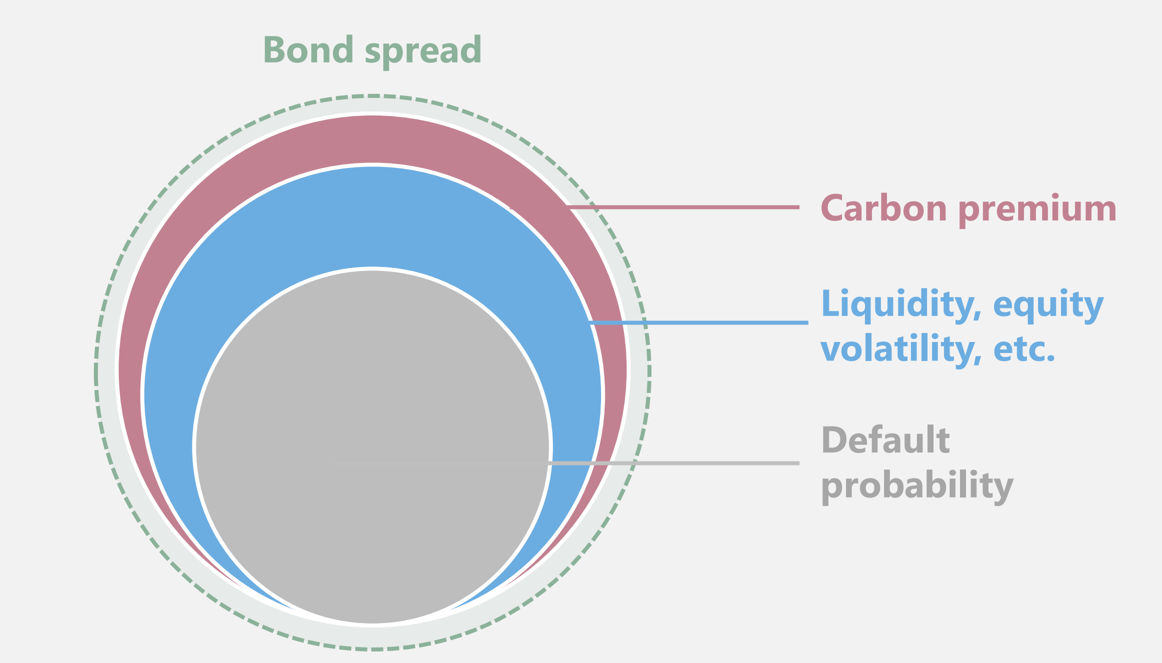

Source: Authors' elaboration.
Diagrammes are intended for illustrative purposes.

Determinants of firm probability of default: illustration

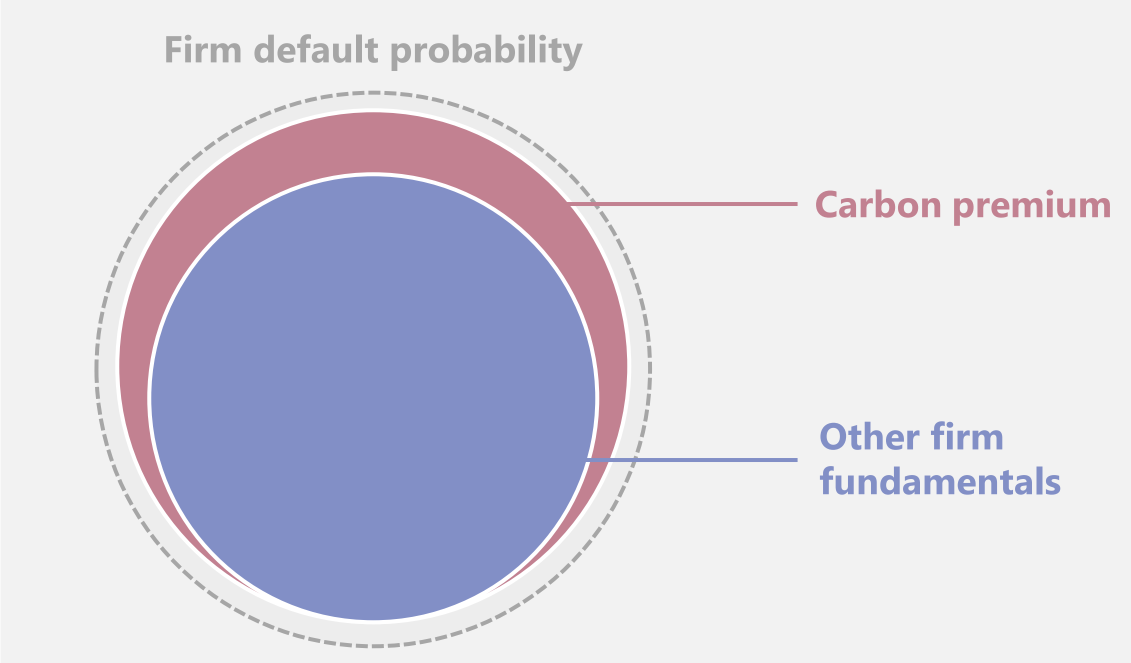

Source: Authors' elaboration.
Diagrammes are intended for illustrative purposes.

Literature review

- Carbon premium in financial assets
  - Bolton and Kacperczyk (2021a, 2021b): firms with higher carbon emissions offer higher returns across sectors and countries
  - Briere and Ramelli (2021): higher investor demand for environmentally responsible stocks is explained by both fundamental and non-fundamental (ie, sentiment) motives. A green sentiment index – which captures shifts in investor appetite for environmental responsibility – has explanatory power on stock price performance.
  - Ehlers et al (2022): there has been a statistically significant carbon premium in lending rates across industries in the syndicated loan market since the Paris Agreement was struck in 2015.
  - Duan et al (2023): bonds from firms with more carbon emissions offer significantly lower returns

- CSR performance/ESG ratings and funding cost:
  - Goss and Roberts (2011); Chava (2014); Ge and Liu (2015); Polbennikov et al (2016); Seltzer et al (2022)

- Determinants of corporate spreads
  - Collin-Dufresne et al (2001); Amato and Remolona (2003); Elton et al (2001); Campbell and Taksler (2002); Chen et al (2007)

1. Measuring carbon footprints

Carbon emissions

- We assume that when making their decisions, investors care about whether the company pollutes the environment or not (carbon emissions), regardless of profit (carbon intensity).
  - This choice is also aligned with regulatory frameworks, such as climate stress tests, which tend to focus on activities with high levels of emissions.

- We consider three types of emissions in our analysis. They follow the Greenhouse Gas Protocol (GGP) classification.
  - **Scope 1**: direct emissions occur from sources that are owned or controlled by the company
  - **Scope 2**: indirect emissions, which come from the firm's purchased electricity, steam and heating/cooling
  - **Scope 3**: other indirect emissions that are a consequence of the activities of a company, but occur from sources not owned by the company. The measure includes both upstream and downstream emissions. We focus on upstream emissions - easier to estimate and have longer time series.
  - While Scope 1 and 2 emissions are part of the disclosure requirement in accordance with the GHG Protocol Corporate Standard; disclosure of Scope 3 emissions are optional.

- We suppose that, when investors look at GHG, they think of them on an **additive basis**. Investors do not consider indirect emissions independent of direct ones. They care about total pollution: the sum of several scopes altogether.

Figure [1.1] Overview of GHG Protocol scopes and emissions across the value chain

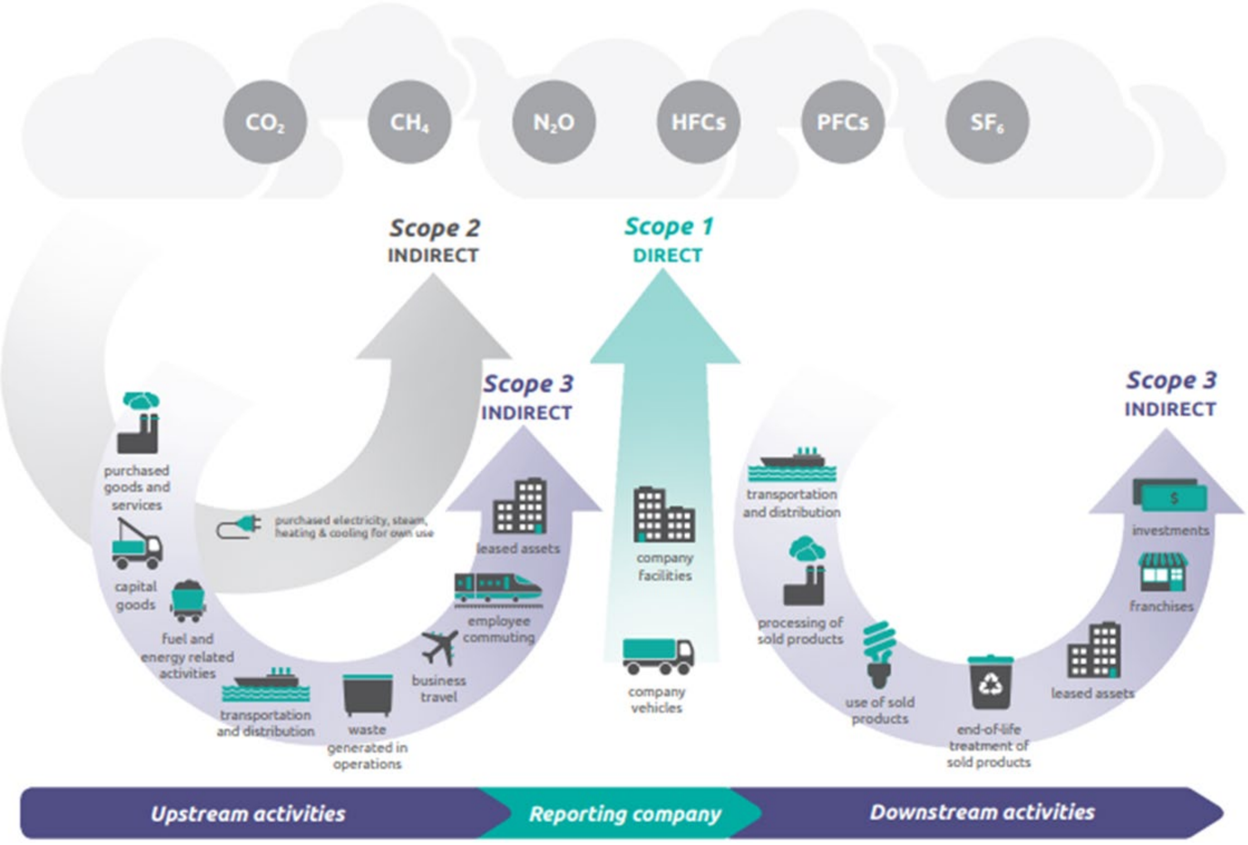

Carbon emissions (cont'd)

- We obtain data on carbon emissions from S&P Global Trucost ("Trucost"), with data since the fiscal year 2002.
- We first focus on the United States, given that its companies are the lion's share of the database in market value terms.
- We choose to start our analysis from 2016, since when the firm composition is more or less stable.

Firm coverage of the S&P Global Trucost emissions database

Number of firms

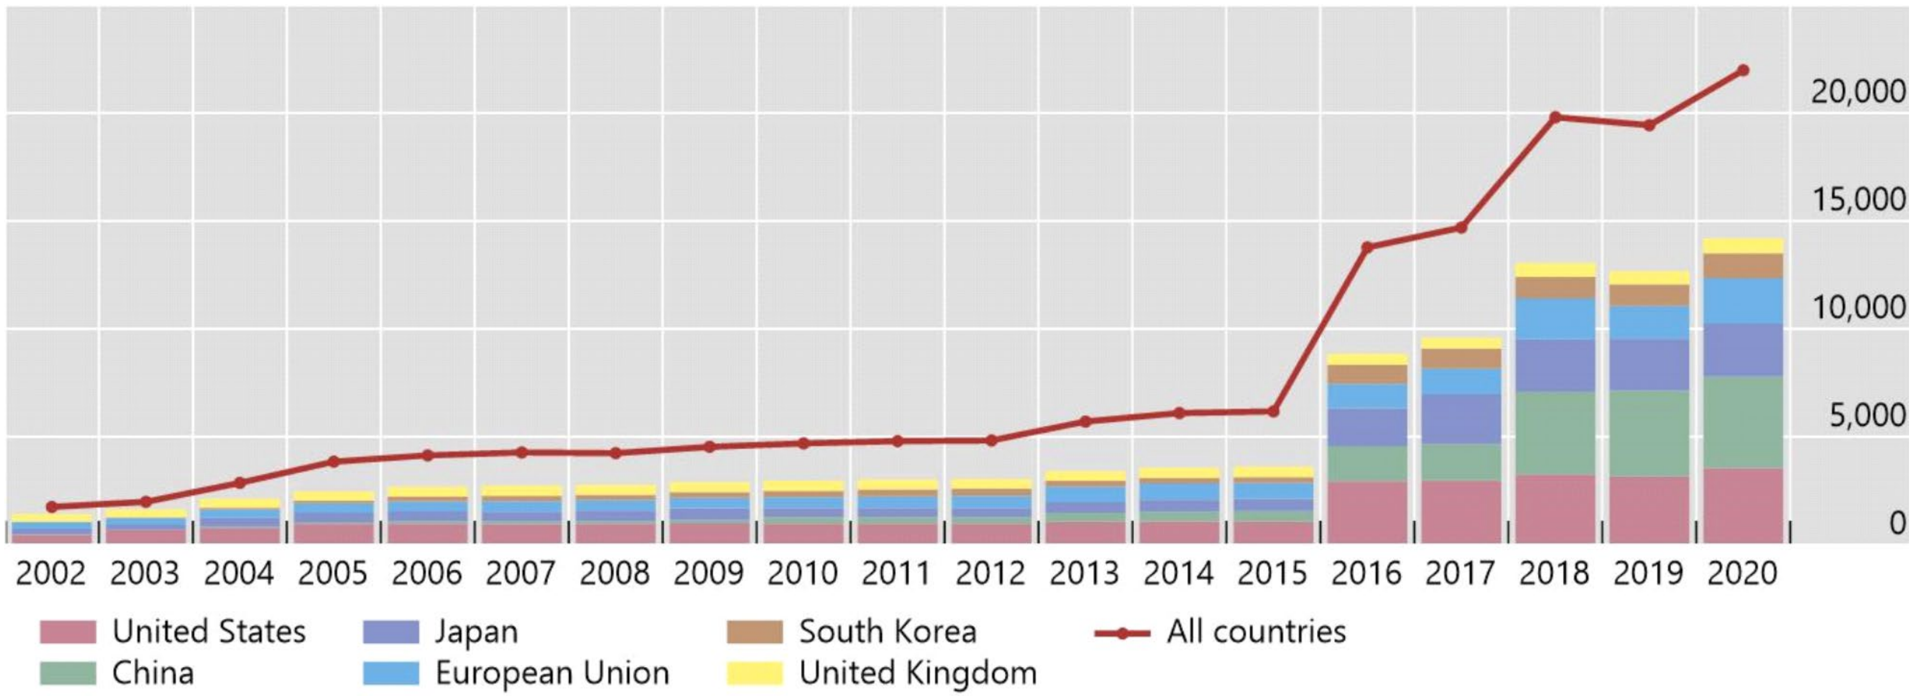

Source: S&P Global Trucost.

Emissions vary by sector

- Across sectors, carbon emissions also display substantial heterogeneity. For direct emissions (red bars), sectors traditionally perceived as brown such as utilities, energy and materials stand out as the top three. When indirect (blue bars) and down-the-value-chain (yellow bars) emissions are included, consumer staples are carbon-intensive, too.
- We call these “energy-intensive” sectors in our analysis, including consumer staples when we talk about 1+2+3.

Average emissions by sector
Thousands of tonnes CO_2

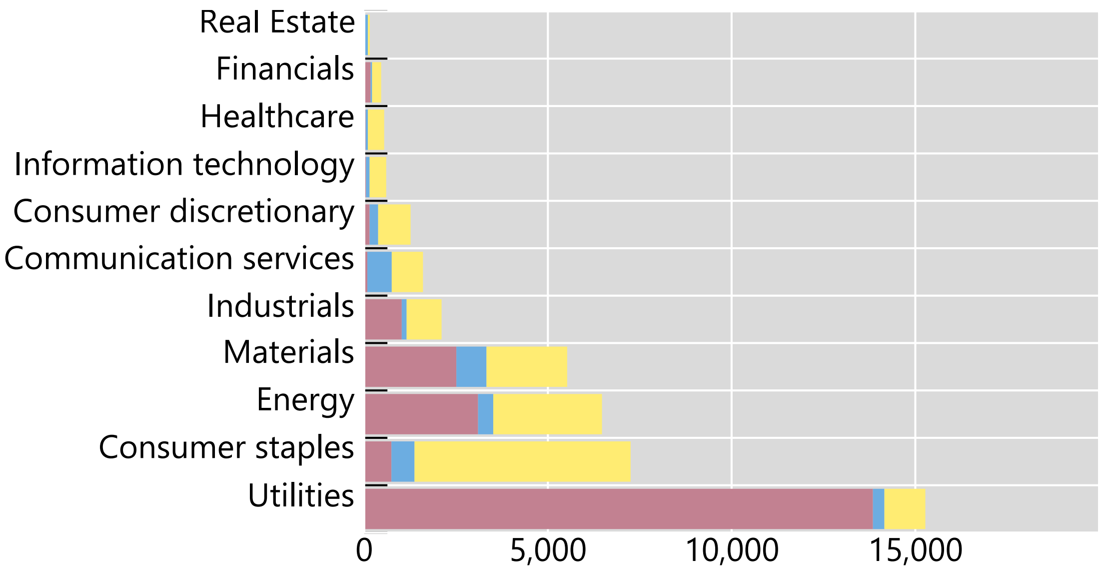

Emissions breakdown by sector
Per cent of total

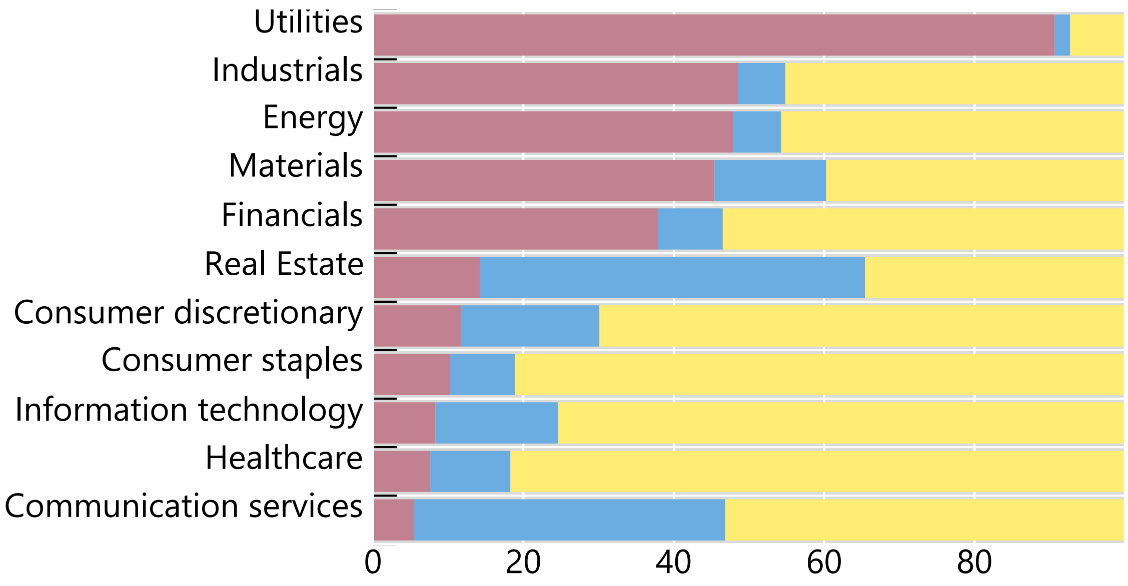

Scope 1 Scope 2 Scope 3

1/ The distribution of carbon emissions is highly skewed. Therefore, a natural logarithm transformation is applied before building the kernels.
Sources: S&P Global Trucost; authors’ calculations.

2. Empirical analysis

2. Empirical analysis

2.1 The preference channel

2.2 The risk channel

Preference channel: the model

- We add a carbon footprint component to a model for spreads that is consistent with literature.[1]

Our exercise: decomposition of option-adjusted spreads
For bond j of firm i at time t

Option adjusted spread[1]	=	Expected default[2]	+	Other known determinants[3]	+	Carbon footprint
$s_{i,j,t}$		$\beta_P P_{i,t}$		$\underline{\beta_Z} Z_{j,t}$		$\boldsymbol{\beta_{P,GHG}} \ln(E)_{j,t-12}$
Corporate bond yield above the risk-free rate.		Captures investors' perceptions of credit risk.		Bond-specific variables such as duration & liquidity.		Firm-level carbon emissions, one proxy at a time.

1/ This decomposition is similar to that of Gilchrist & Zakrajsek (2012) and is motivated in part by the existence of the "credit spread puzzle," the well-known result from the corporate finance literature, showing that less than one-half of the variation in corporate bond credit spreads can be attributed to the financial health of the issuer.
2/ Bloomberg 5-year ahead firm-level default probability (physical measure).
3/ Determinants as registered in the literature by Elton (2001), Campbell & Taklser (2002) and Chen, Lesmond & Wei (2007). They include bond-specific controls and month-year fixed effects.

The effects of carbon emissions on risk-adjusted OAS

Effect of carbon emissions on corporate spreads[1]
Coefficient on lagged *ln*(emissions)

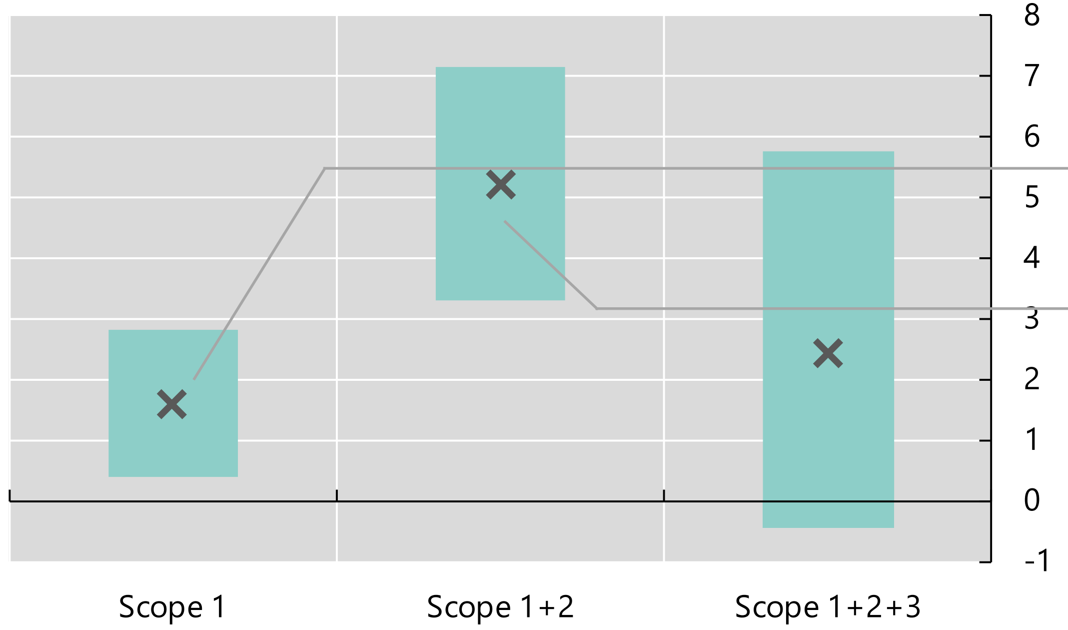

1/ Bars denote confidence intervals at the 90% level. Interpretation: if the bar touches zero, the coefficient does not pass the test at the 10% level.
Specifications with time-, firm- and credit rating fixed effects. Standard errors clustered at the security level.
Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

An increase in direct carbon emissions leads to an increase in corporate spreads via $\beta_{P,GHG}$.

- Concretely, a 1% increase in CO_2 emitted directly entails a 0.02 basis point yield increment.
- A 1% increase in indirect CO2 entails a 0.05 basis point increment.
- The coefficient for scope 1+2+3 is positive but *not* statistically significant (mind the scope 3 data!).

The effects of carbon emissions on risk-adjusted OAS (cont'd)

- These specifications show are evidence of a **statistically significant** carbon risk premium.
- We can also look at a positive message: what happens when emissions are reduced?
- For instance, what happens to funding costs when companies halve their direct carbon emissions? Computing what happens when emissions change by -50%, the effect appears **economically low:** 1 to 4 basis points.

Decrease in spread implied by a 50% reduction in emissions

Spread change in basis points

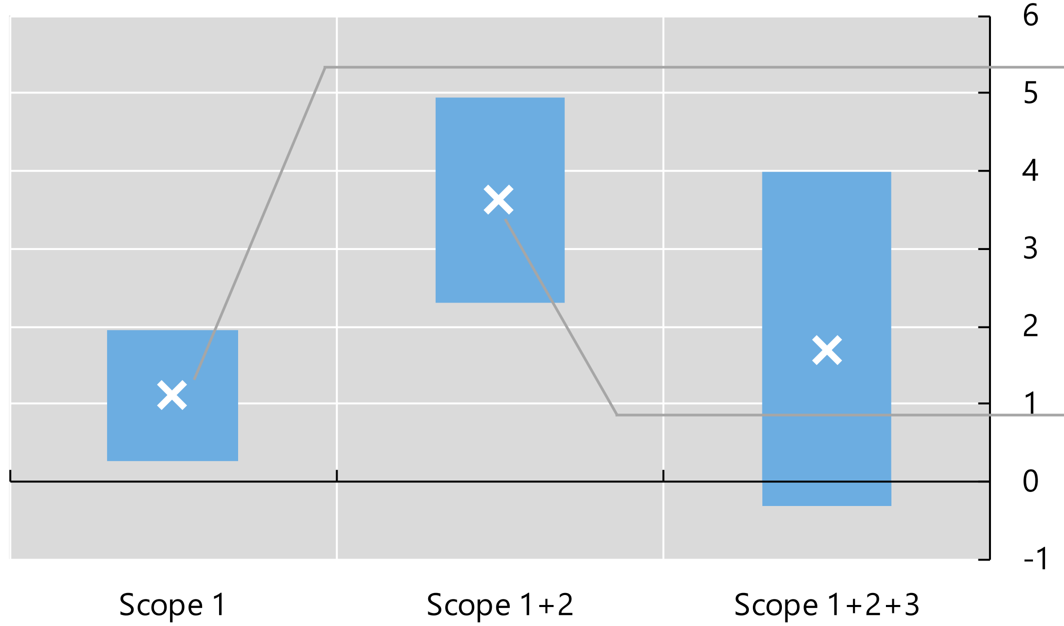

A 50% decrease in total emissions implies a 1.1 basis point drop in the bond spread.

The effect is slightly higher at 4 basis points when scopes 1 and 2 and considered jointly.

1/ Bars denote confidence intervals at the 90% level. Interpretation: if the bar touches zero, the coefficient does not pass the test at the 10% level. Specifications with time-, firm- and credit rating fixed effects. Standard errors clustered at the security level.

Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

Looking by sector

- There are important differences in average emissions between companies considered "**energy intensive**" and those that are not. So, should bonds from all firms bear the same carbon premium? We address this question here.
- We interact an energy-intensive indicator variable with total carbon emissions and re-run our analysis.
- The impact within the set of energy-intensive firms is of **non-negligible economic importance**.

Effect of carbon emissions on corporate spreads[1]
Coefficient on lagged *ln*(emissions) × energy intensive dummy[2]

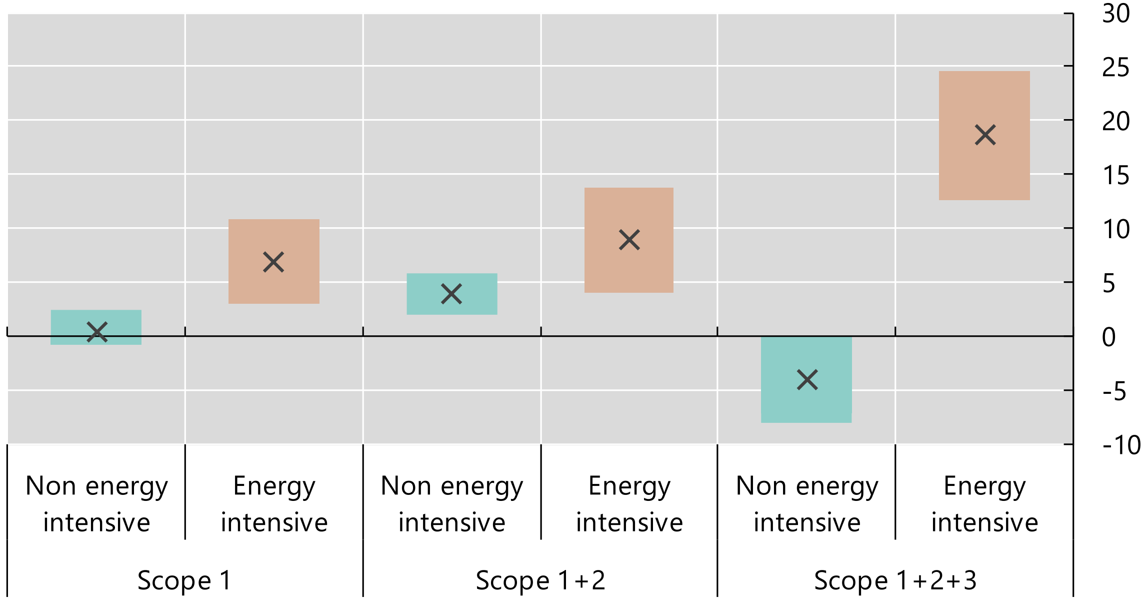

Decrease in spread implied by a 50% reduction in emissions
Spread change in basis points

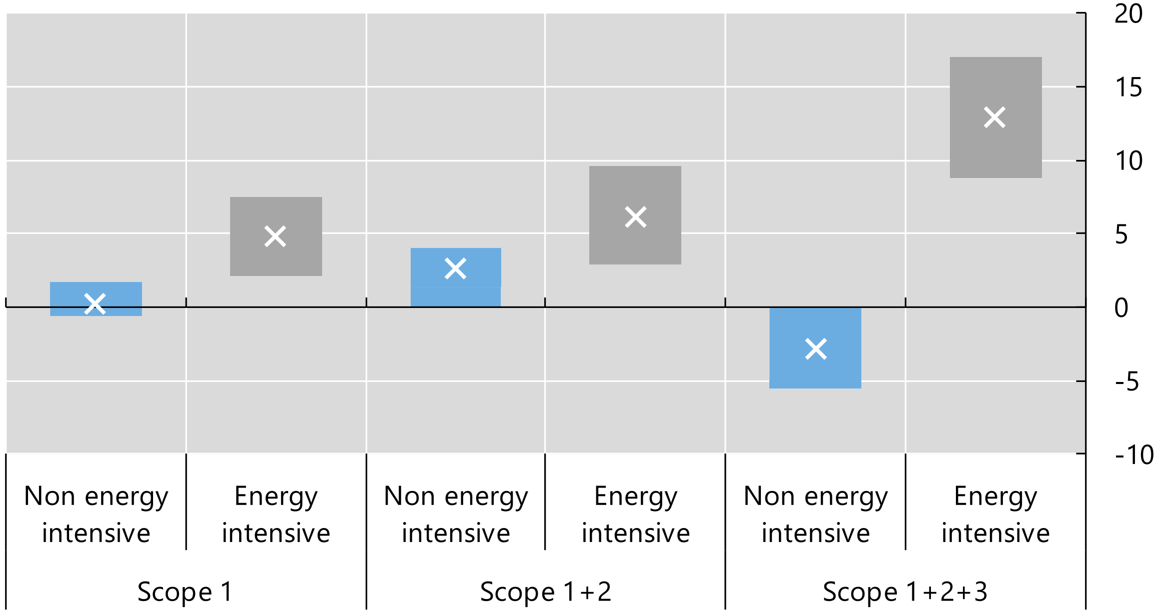

1/ Bars denote confidence intervals at the 90% level. Interpretation: if the bar touches zero, the coefficient does not pass the test at the 10% level.
2/ The value of the dummy is equal to 1 when the company belongs in the category and zero otherwise.
Specifications with time-, firm- and credit rating fixed effects. Standard errors clustered at the security level.
Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

Looking by maturity

- Next, we ask ourselves whether carbon risk compensation may also be related to the term of each instrument. The analysis is possible thanks to the security-level approach we have taken in our models.
- To do so, we create a dummy variable within buckets, using five-year steps.
- For scope-1 emissions, we find statistical significance across most maturity buckets, except the shortest (<5 years).

Effect of carbon emissions on corporate spreads[1]

Coefficient on *ln*(scope 1 emissions) × maturity dummy[2]

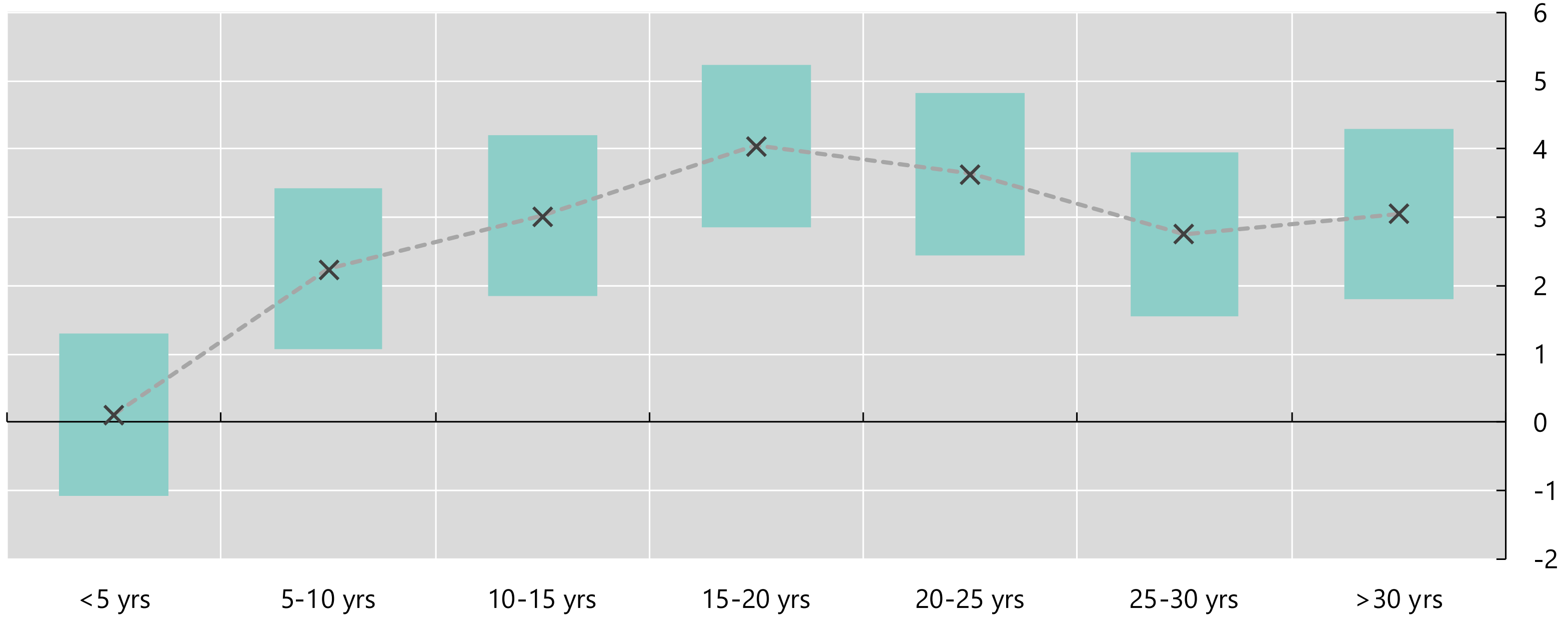

1/ Bars denote confidence intervals at the 90% level. Interpretation: if the bar touches zero, the coefficient does not pass the test at the 10% level. 2/ The first bucket is less than five years in maturity, then with bonds between 5 and 10 years, and so forth – our last bucket comprises all bonds above 30 years. Specifications with time-, firm- and credit rating fixed effects. Standard errors clustered at the security level. Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

Putting it all together

- We have found that both sector and maturity matter. What happens if we interact both effects at the same time?
- We uncover that the maturity loading structure is statistically significant for both cases: energy-intensive (EI) and non-energy-intensive (NEI) firms. And their order of magnitude is different.
- Below, we showcase results for scope 1 + 2 emissions together, which we consider our key finding of this section.

Effect of carbon emissions on spreads[1]

Coefficient on *ln*(scope 1+2 emissions) × maturity × sector[2]

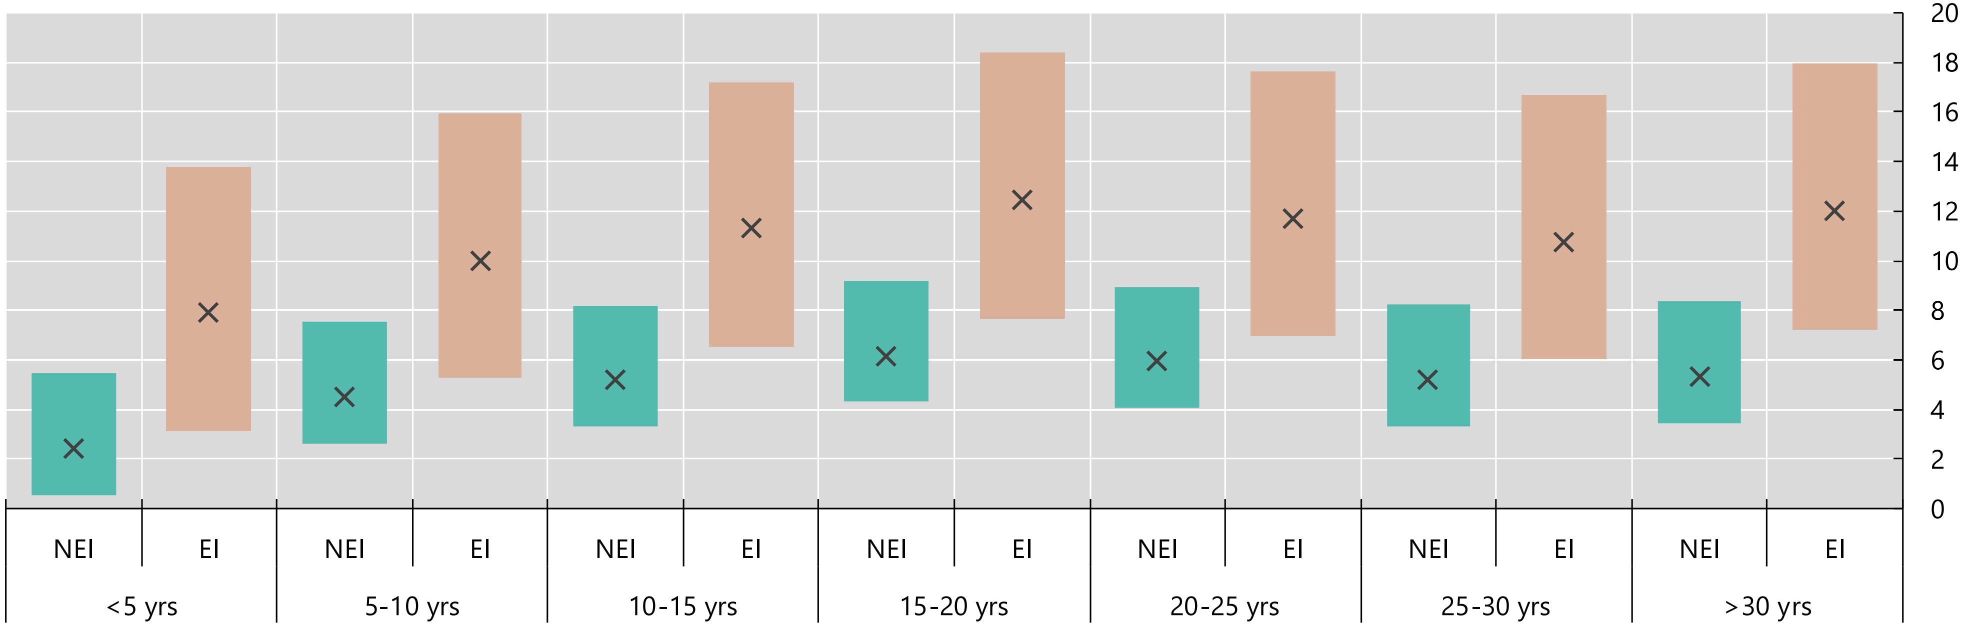

NEI = Non-energy intensive sectors; EI = energy-intensive sectors.
1/ Bars denote confidence intervals at the 90% level. Interpretation: if the bar touches zero, the coefficient does not pass the test at the 10% level.
2/ Combines the energy-intensive dummy with a maturity bucket dummy. See slides 18 and 19 for individual details. Standard errors clustered at the security level.
Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

The term structure of carbon risk premia

- Our findings allow us to speak about a term structure or "curve" of carbon premia (plural) which is hump-shaped. Why?
  - We offer two theories: (1) the long term-nature of environmental risks; and (2) preferred habitat.
- The differentiated impact across industries probably reflects investors' greater scrutiny of firms traditionally viewed as brown. The effect for energy-intensive firms can be equivalent to a credit-rating upgrade of 0.8 notches.[1]

Decrease in spread implied by a 50% reduction in scope 1+2 emissions

Spread change in basis points

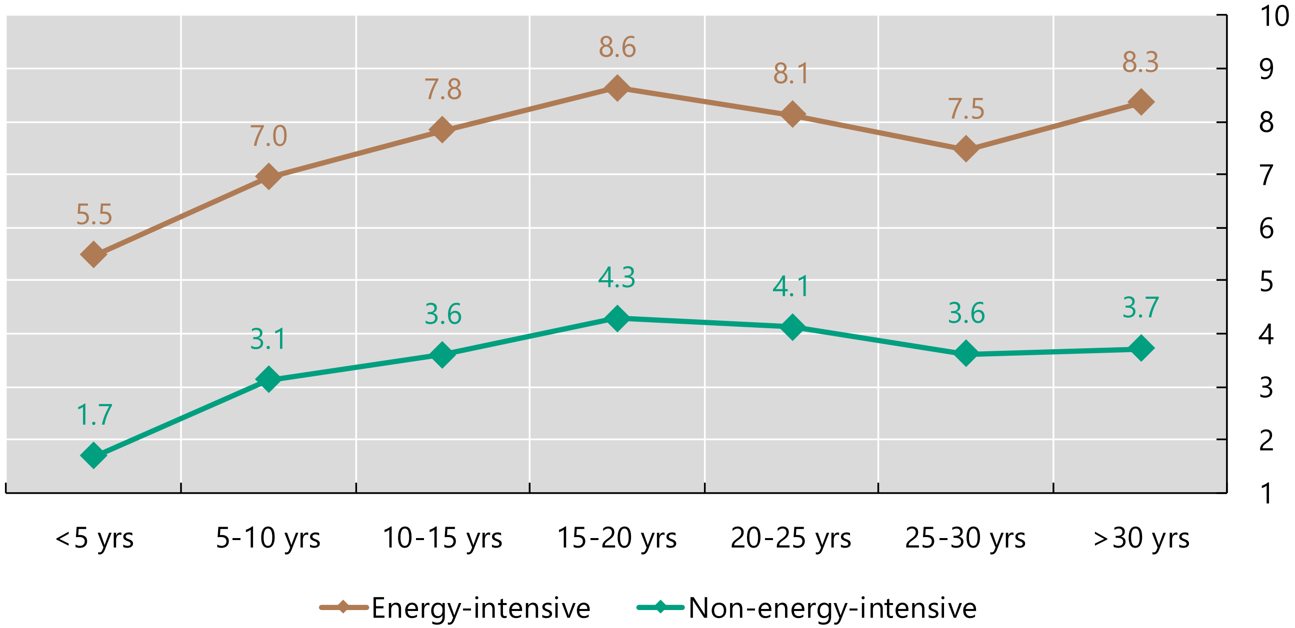

Sources: BIS calculations; Bloomberg, Refinitiv, Trucost.

1/ This result is the average of individual notch changes across all bonds. To estimate each individual notch change, the 8.6 basis point impact from a 50% change in emissions is divided by the differential between two (mean) spreads: that of the bond's credit rating and that of the adjacent notch below. Our sample is comprised of bonds with ratings ranging from AAA to C on the Fitch scale.

2. Empirical analysis

2.1 The preference channel

2.2 The risk channel

Risk channel: the model

- Our model consists of decomposing firm-level default probabilities into two components: one related to the financial health of the firm, and another related to carbon emissions.
- However, given that probabilities – our regressand – are bounded between 0 and 1, we perform a logit transformation first. We re-express firm-level annualised default probability as $P_{i,t}$ as the log-odds of default $\tilde{P}_{i,t}$.[1]

Another exercise: decomposition of default probabilities
For bond i of firm j at time t

Odds of default[1]		Known determinants[2]		Carbon footprint		
$\tilde{P}_{i,t}$	=	$\underline{\beta_X} X_{j,t}$	+	$\boldsymbol{\beta_{R,GHG}} \ln(E)_{j,t-12}$	+	$\epsilon_{j,t}$
Investors' perceptions of credit risk.		Firm-specific fundamentals and fixed effects.		Firm-level carbon emissions, one proxy at a time.		

1/ $\tilde{P}_{i,t}$ is the logit transformation of such probability for firm i at time t calculated as $\log(p_{i,t}/1 - p_{i,t})$. Also, this way, our specification is consistent with commonly used logit model that models default events.
2/ $X_{i,t}$ is a set of firm-level control variables, including size of assets, long-term debt/assets ratio, earnings/asset ratio, capital/assets ratio and return on assets. Includes a vector of time and sector fixed effects.

Risk channel: baseline estimation results

- For all the models, coefficients in front of control variables are of right signs and statistically significant.
- Zooming into models that consider carbon emissions, our result confirms our hypothesis – a firm's carbon emissions have a positive correlation with their probability of default ("PD"). For all carbon emission measures, $\beta_{R,GHG}$ is above zero and statistically significant.

Effect of carbon emissions on default probability[1,2]
Coefficient on lagged *ln*(emissions)

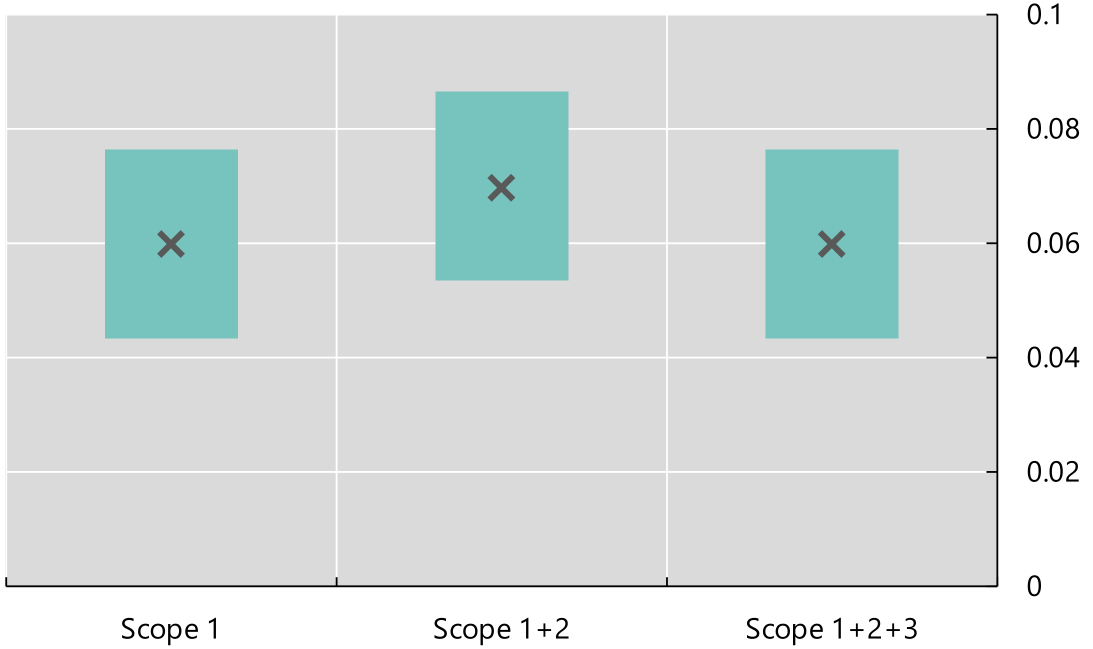

1/ Bars denote a 90% confidence interval. All coefficients are statistically significant at the 1% level. Specifications with time- and sector-fixed effects. Standard errors clustered at the firm level.
2/ The dependent variable is the logit transformation of such probability for firm i at time t calculated as $\log(p_{i,t}/\ 1 - p_{i,t})$.
Sources: BIS calculations; S&P Capital IQ, S&P Global Trucost.

Risk channel: sector effects

- As in the preference channel, we test whether the impact is only viable in certain sectors – for example, those commonly viewed as "brown" industries. Our results suggest that the impact prevails regardless of energy intensiveness.
- For scope 1 and scope 1 and 2 emissions jointly, the impact on probabilities of default is larger in energy-intensive sectors.
- These results are for the odds of default. How to interpret these results in terms of default probability?

Effect of carbon emissions on default probability[1]
Coefficient on lagged *ln*(emissions) × energy intensive dummy[2]

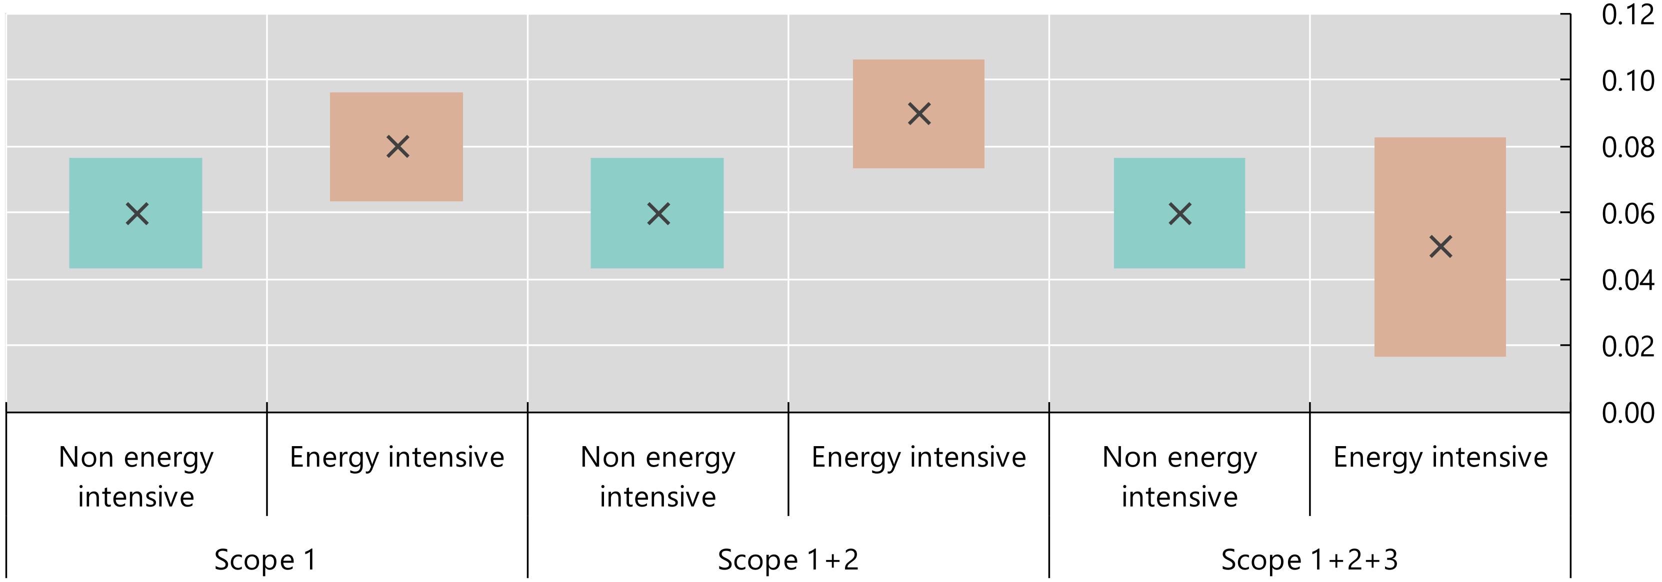

1/ Bars denote a 90% confidence interval. All coefficients are statistically significant at the 1% level. Specifications with time- and sector-fixed effects. Standard errors clustered at the firm level.
2/ The value of the dummy is equal to 1 when the company belongs in the category and zero otherwise.
Sources: BIS calculations; S&P Capital IQ, S&P Global Trucost.

Impact on default probability

- What is the effect on corporate spreads? Since the **transformation is non-linear** from $P_{i,t}$ to $\tilde{P}_{i,t}$, the impact of carbon emissions on probability of default – and thus – spreads is not a constant.
- We consider default probabilities up to 32% to compute the **impact for each PD level**. For example, for firms with a 1% probability of default, the change impact a 0.04 pp increase in PD for a 50% reduction in total GHG emissions.

Distribution of firm-level physical probabilities of default[1]

Probability density

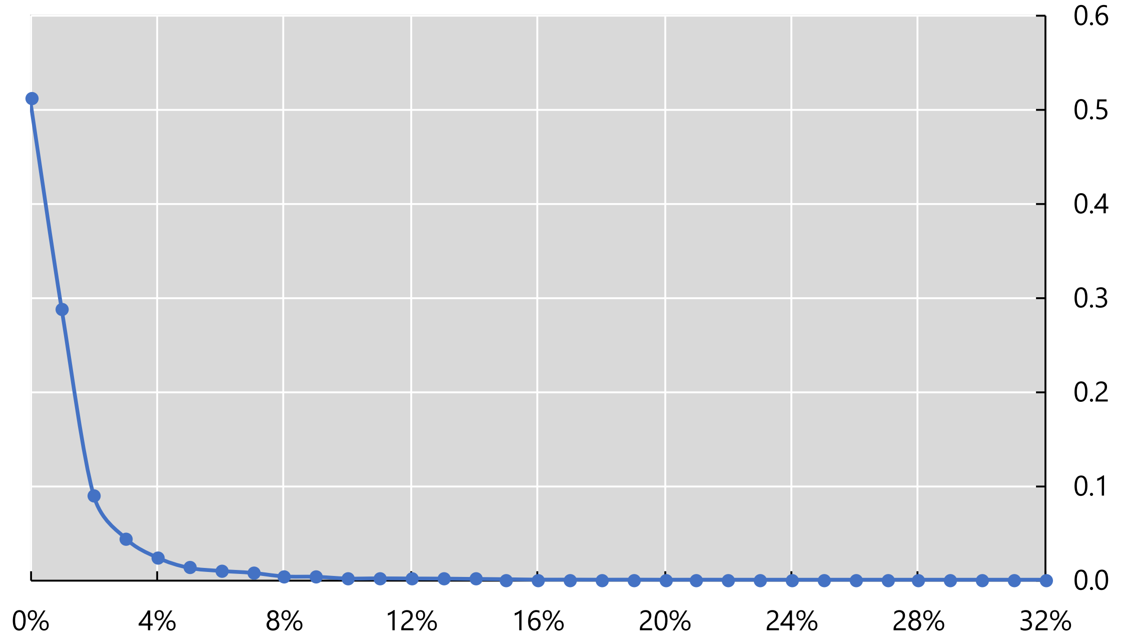

Decrease in PD due to a 50% decrease in emissions[2]

PD change in percentage points

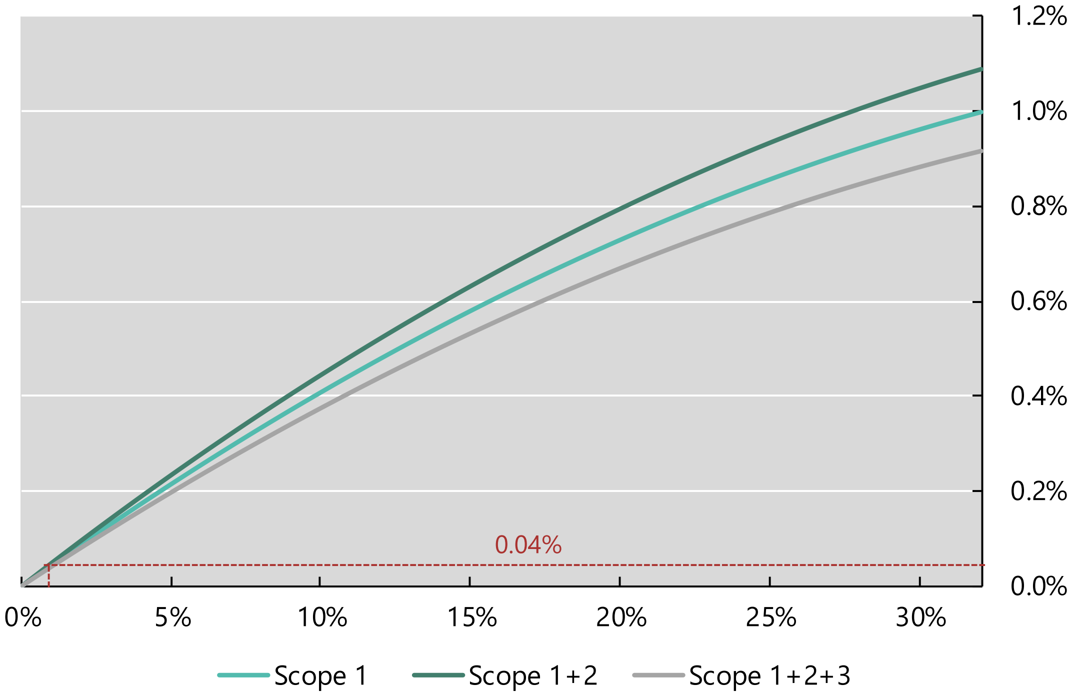

Source: BIS calculations using Refinitiv, S&P Trucost and Bloomberg data.
1/ Annualised, 5-year ahead physical default probability.
2/ Computed for each probability level from 0% to 32% as the inverse logit transformation of: the change in log-odds of default implied by the change in log-scope 1 emissions (%).

Risk channel impact on corporate spreads

- By multiplying our estimated change in $P_{i,t}$ by the coefficient on β_P in our preference channel model (the first term of the equation), we can find the total impact of the risk channel on spreads. For example: -50% in scopes 1+2.
  - For a firm with 1% PD, the impact is of **1.3 bp for NEI firms** and **2.2 bp for EI firms**.

Decrease in spread implied by a 50% reduction in emissions: via the risk channel, for non-energy-intensive firms[1]
Spread change in basis points

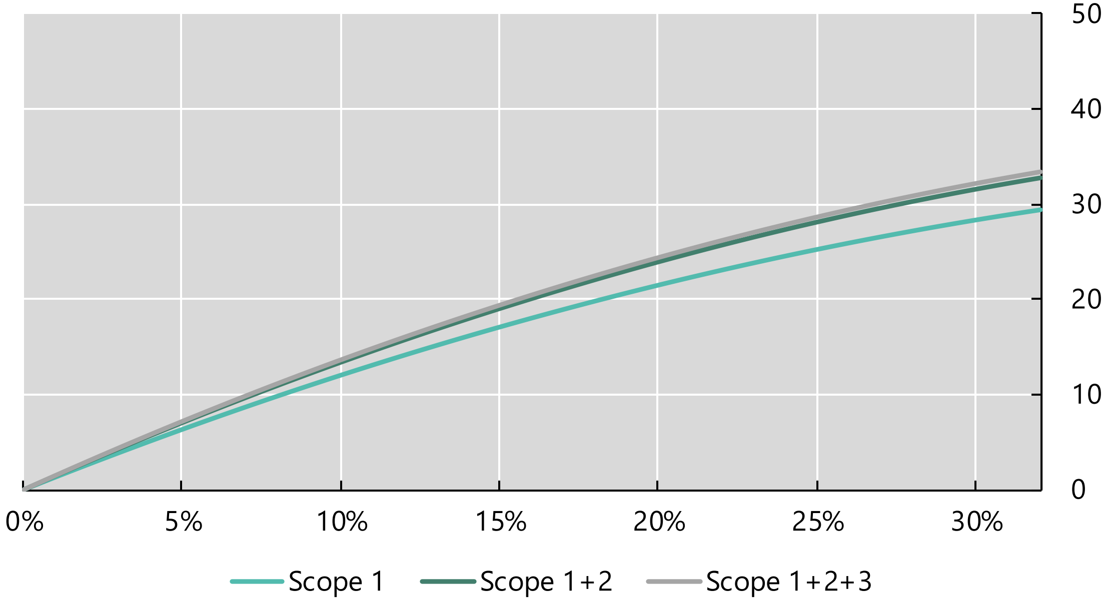

Decrease in spread implied by a 50% reduction in emissions: via the risk channel, for energy-intensive firms[1]
Spread change in basis points

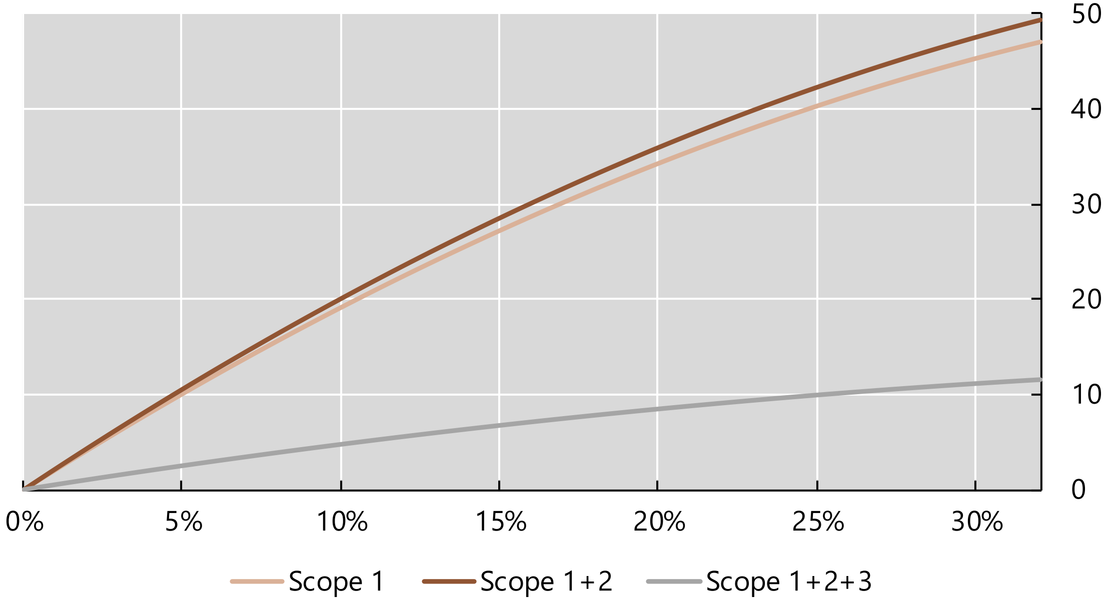

Source: BIS calculations using Refinitiv, S&P Trucost and Bloomberg data.
1/ The graphs display the estimates of a change in option adjusted spread using the risk channel for each level of default probability from 0 to 32%. These are then multiplied by the coefficient of PD in the preference channel equation.

Putting both channels together: the term structure revisited

- To wrap up, we can compute both effects together to get our final term structure of carbon risk premia. The preference channel reflects a maturity structure, while the risk channel may dominate depending on how close a firm is to default.

Decrease in spread implied by a 50% reduction in emissions, base probability of default = 1%[1]

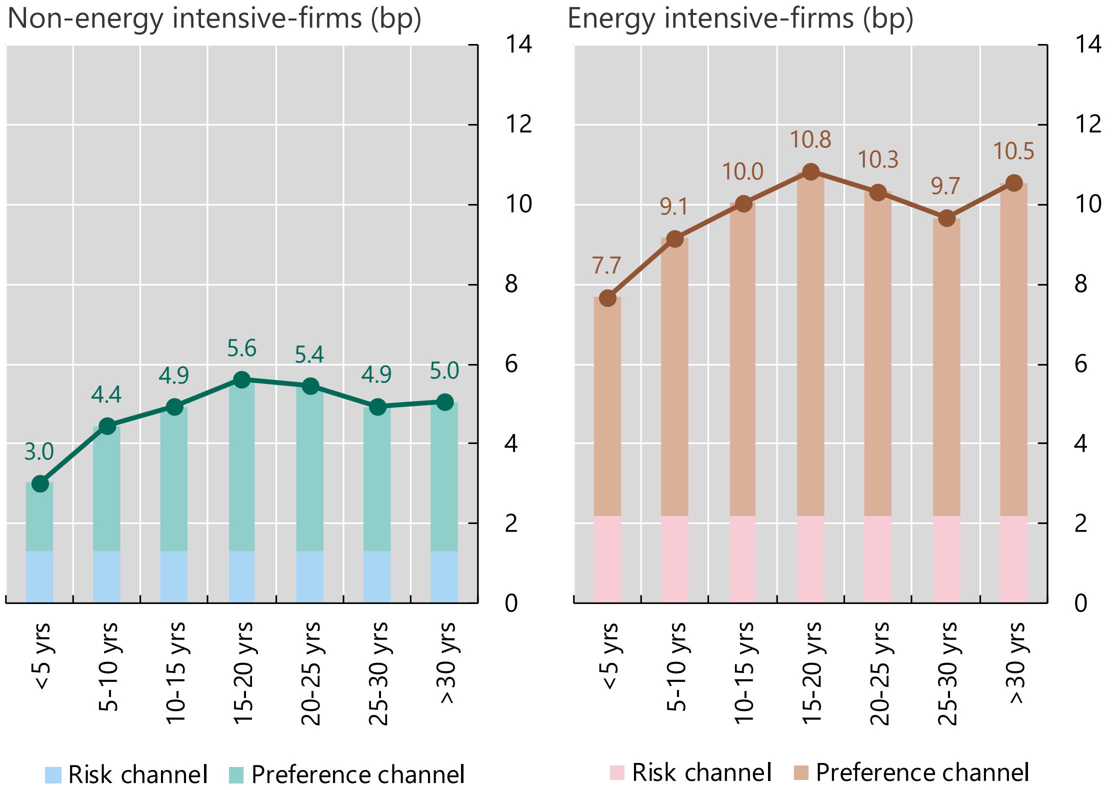

Decrease in spread implied by a 50% reduction in emissions, base probability of default = 10%[1]

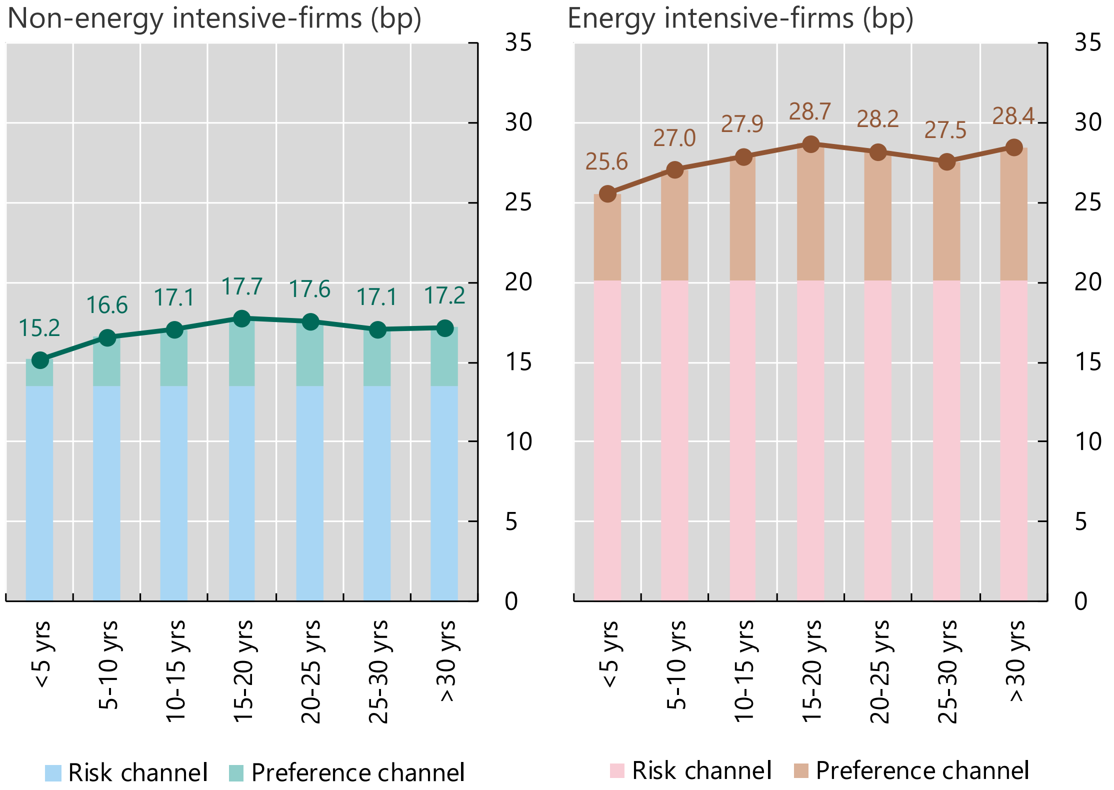

Source: BIS calculations using Refinitiv, S&P Trucost and Bloomberg data.

1/ The graphs adds up the effects of a 50% decrease in emissions via both the preference and risk channels.

Robustness check and additional analysis

- Preference channel
  - Different measure of PD and liquidity
- Risk channel
  - Different measure of PD
- Total carbon premium
  - Estimation in one go without decomposing different channels
- Premium for carbon intensity

3. Concluding remarks

Conclusions

For the US, we find a statistically significant carbon premium via both the preference and risk channels.

Preference channel

Bonds from firms with higher carbon emissions tend to trade at marginally higher risk-adjusted yields.

Risk channel

Carbon emissions also matter for firm default – higher emissions increase its likelihood.

Stylised facts

- The **size** of the carbon *premia* is larger for firms that are energy-intensive.
  - Under the preference channel, it can reach almost a notch change in credit rating terms, if firms aim to halve their emissions.
  - Under the risk channel, and for the average firm (1% PD), it is a couple basis points.

- The carbon premia is **composed,** mainly of preference effects.
  - There is a maturity structure, as it appears ESG-aware investors operate in specific sectors.
  - However, the premia is quickly dominated by credit risk effects for those firms with high(er) credit risk.

Implications

Firms need to be aware that their funding costs may rise should they increase their GHG. Conversely, they may enjoy the benefits of cheaper funding costs.